

METHOD FOR PROVIDING LONG-LIVED BROADCAST ENCRYPTION

BACKGROUND

5 1. Technical Field:

The present application relates generally to broadcast encryption and, more particularly, to a long-lived broadcast encryption scheme that adapts to the presence of pirate decoders and maintains the security of broadcast to authorized users as encryption keys are compromised over time.

10 2. Description of Related Art:

In general, broadcast encryption (BE) techniques are employed to encrypt digital content to ensure that only privileged users are able to recover the content from an encrypted broadcast. Keys are allocated in such a way that users may be prevented on a short-term basis from recovering the message from the encrypted content. This short-term exclusion of users occurs, for example, when a proper subset of users request to view a movie. The long-term exclusion (or, revocation) of a user is necessary when a user leaves the system entirely.

In practice, broadcast encryption schemes are typically smartcard-based, wherein key material is held in a “tamper-resistant”, replaceable smartcard. These smartcards, 20 however, may be used to construct pirate smartcards (or pirate decoders) that allow non-paying customers to recover content. For instance, a coalition of unscrupulous users may conspire to attack a BE system by breaking open their smartcards to extract the keys and build pirate decoders using the extracted decryption keys, allowing non-authorized,

nonpaying users to utilize the pirate decoders to recover the content of encrypted broadcasts.

There are various conventional methods that have been employed to guard against the use of such pirate decoders. For instance, a broadcast encryption scheme may be
5 coupled with a traceability (or traitor tracing) scheme to offer some protection against piracy. The traceability scheme is employed to render the practice of building pirate smartcards risky. This is accomplished by allocating keys to users in such a way that, once such a pirate smartcard is confiscated, at least one of the cards that was used to construct it can be identified.

10 There are disadvantages associated with such conventional schemes. For instance, traitor tracing schemes are typically designed having x -traceability, wherein it is possible to identify at least one of the smartcards used to construct a given pirate card only if x smartcards, at most, are used to construct the pirate card. Therefore, the security achieved in traceability schemes is limited by the necessity to have a bound on
15 the number of users in a coalition. Another disadvantage associated with traceability schemes, as well as other conventional broadcast encryption schemes, is that they do not take into account the effect of the compromised keys when encrypting the content. Indeed, conventional broadcast encryption schemes do not provide a mechanism for providing continued, secured broadcasting to privileged sets of users after compromised
20 keys that are contained on pirate decoders and/or smartcards of revoked users (whose contract has expired) have been identified. Instead, some schemes require that a private communication be made to each of the remaining users (e.g., over a network) when a single user is revoked in order to maintain the ability to revoke (or exclude) a certain

amount of users. In a BE scheme employing smartcards, however, the cost of reprogramming or replacing a large number of cards each time a user is revoked or a pirate smartcard is identified is extremely costly.

Accordingly, a broadcast encryption scheme that is not dependent on the size of
5 the coalition constructing a pirate decoder and that retains an ability to broadcast securely
in the presence of pirate smartcards, and consequently, compromised keys, is highly
desirable.

SUMMARY OF THE INVENTION

The present invention is directed to a long-lived broadcast encryption method that
10 adapts to the presence of compromised keys and continues to broadcast securely to
privileged sets of users over time. In one aspect of the invention, a method for providing
long-lived broadcast encryption comprises the steps of: allocating, to each of a plurality
of subscribers, a corresponding set of subscriber keys; broadcasting encrypted content to
the plurality of subscribers using a set of broadcast keys, wherein the encrypted content is
15 decoded by a given subscriber using the subscriber's corresponding set of subscriber
keys; modifying the set of broadcast keys, which are used for broadcasting encrypted
content, by excluding compromised subscriber keys; and updating a set of subscriber
keys corresponding to at least one subscriber when the at least one subscriber's set of
subscriber keys comprises an amount of active keys that falls below a first predetermined
20 threshold.

In another aspect of the invention, the set of subscriber keys allocated to a given
user is encoded on a smartcard which is issued to the subscriber. When a pirate decoder
is discovered, it is analyzed and the keys it contains are identified. Such keys are called

“compromised,” and are not used henceforth for broadcasting. Similarly, when a subscriber’s contract expires, the subscriber is “excluded,” and each of the keys contained on the excluded subscriber’s smartcard are considered compromised. The compromised keys are excluded from the broadcast encryption process. Over the course of time, while

5 compromised smartcards are identified, a state may exist in which the number of compromised smartcards (and consequently, the number of compromised keys) reaches a threshold above which secure communication to authorized users is not possible (because the number of active keys on their smartcards fall below a predefined threshold amount). To restore secure broadcast capabilities to such users, their smartcards are replaced with a

10 new smartcard comprising a new set of keys.

In yet another aspect of the invention, a long-live broadcast encryption scheme is reactive, being driven by the card compromises: either due to pirate decoders or the expiration of users’ contracts. These card compromises are used to divide time into administrative “epochs” comprising d compromises each. At the end of an epoch, the

15 service provider determines which authorized subscribers require card replacement, and such cards are replaced. Therefore, the primary cost in a long-lived BE scheme according to the present invention is the amount of rekeying that is required in each epoch.

In another aspect of the present invention, a long-lived broadcast encryption

20 scheme employs an efficient (short term) randomized broadcast encryption scheme that yields an (m, α) - cover-free family to ensure with high probability that any coalition m smartcards can not cover a fraction α , of the keys contained on another smartcard. A $(\alpha r, |S_P|)$ - threshold protocol is used to broadcast to a privileged set P (authorized

subscribers) using a set of keys S_P , such that a fraction α of a total number of keys r contained on the smartcard of an authorized user is needed to recover the encrypted broadcast content. To render a compromised card useless as a decoder, all the keys contained on the compromised card are excluded from S_P , creating a new set S_p^1 .

5 Thereafter, broadcasting to the privileged set P is performed with an $(ar, |S_p^1|)$ - threshold protocol. A recording session is entered whenever d cards become unavailable. In a recording session, new values are chosen randomly for all compromised keys, and then any user with less than ar active keys receives a new card. The parameter d in the fixed schedule will be based on the number of compromised users and the desired
10 transmission length.

In yet another aspect of the invention, one or more parameters of a long-lived broadcast encryption scheme can be selected, a priori, to ensure that only a bounded number of users need recording in each epoch

These and other aspects, features, and advantages of the present invention will
15 become apparent from the following detailed description of the preferred embodiments, which is to be read in connection with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of a long-lived broadcast encryption system according to
20 an embodiment of the present invention;

Fig. 2 is a flow diagram of a method for providing long-lived broadcast encryption according to one aspect of the present invention;

Fig. 3 is a diagram illustrating different periods (epochs) of a randomized broadcast encryption scheme according to one aspect of the present invention;

Fig. 4 is a graphical diagram illustrating experimental results of smartcard replacement based on different values of the total number of keys in a long-lived broadcast encryption scheme according to one aspect of the present invention;

Fig. 5 is a graphical diagram illustrating experimental results of smartcard replacement based on different values of the number of keys per smartcard in a long-lived broadcast encryption scheme according to one aspect of the present invention; and

Fig. 6 is a graphical diagram illustrating experimental results of smartcard replacement based on different values of epoch length in a long-lived broadcast encryption scheme according to one aspect of the present invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

It is to be understood that the system modules and method steps described herein may be implemented in various forms of hardware, software, firmware, special purpose processors, or a combination thereof. In particular, the present invention is preferably implemented as an application comprising program instructions that are tangibly embodied on a program storage device (e.g., magnetic floppy disk, RAM, ROM, etc.) and executable by any device or machine comprising suitable architecture. It is to be further understood that, because some of the constituent system modules and process steps depicted in the accompanying Figures are preferably implemented in software, the actual connections between such components may differ depending upon the manner in which the present invention is programmed. Given the teachings herein, one of ordinary skill in

the related art will be able to contemplate these and similar implementations or configurations of the present invention.

Referring now to Fig. 1, a block diagram illustrates a system for providing long-lived BE (broadcast encryption) according to an embodiment of the present invention. In general, a long-lived BE system provides a mechanism for adapting a broadcast encryption scheme to the presence of compromised broadcast encryption keys that are contained on compromised (pirate) smartcards and smartcards of a subscribing user whose membership has expired or has been terminated ("excluded user"). The exemplary system 10 of Fig. 1 comprises a trusted server 11 associated with a service provider of certain broadcast content 12. In general, the trusted server 11 maintains security by authenticating and authorizing subscribers, and transmits encrypted broadcasts to privileged users (authorized subscribers) over a network 13 (e.g., LAN, WAN, Internet, etc). In a preferred embodiment, the service provider issues each subscriber a smartcard S_1, \dots, S_n comprising a set of decryption keys that are used for decoding encrypted broadcasts transmitted from the content provider (e.g., inserting the smartcard in a reader of a set-top box to decrypt a broadcasted event).

In the following discussion, the set of all subscribers/users (privileged and non-privileged/excluded) is denoted $\{u_1, \dots, u_n\}$, for $i = 1 \dots n$ where subscript n represents a total amount of users. The set of keys associated with the i^{th} user is denoted by U_i .

Furthermore, a privileged sets of users, denoted P , is of a fixed size $n - m$, where n represents the set of all users (as indicated above) and m represents a number of users that are non-privileged/excluded (a set of m excluded/non-privileged users is denoted as X).

The trusted server 11 comprises a broadcast encryption module 14 and a key manager/generator module 15, which together, provide a mechanism for long-lived BE according to the present invention. In general, the broadcast encryption module 14 employs a broadcast encryption technique based on symmetric key (or private key) 5 encryption, and a broadcasting protocol for securely distributing broadcast content to a set of privileged users. More specifically, in a preferred embodiment, the broadcast encryption module 14 employs (1) an efficient (short-term) BE scheme (such as a randomized BE scheme described below) that yields an (m, α) -cover free family, and a unicast key between each user and the center; and (2) an $(s, |S_p|)$ - threshold protocol to 10 broadcast to a privileged set of users $P = \{u_1 \dots u_{n-m}\}$, where S_p denotes a set of keys used to broadcast to the privileged set P . If some user is unreachable under the $(s, |S_p|)$ - threshold protocol (i.e., as explained below, if too many of the user's keys appear on "dead" cards), then the unicast key will be used to reach that user. It is to be understood 15 that a long-lived BE scheme according to a preferred embodiment of the present invention is an extension of the randomized BE scheme described, for example, in "Long-Lived Broadcast Encryption," by J. Garay, et al., Advances in Cryptology - Crypto '00, Lecture Notes in Computer Science 1880 (2000), pp. 333-352. It is to be understood, however, that (as discussed below) other BE schemes may be employed herein, preferably, BE schemes that are based on cover-free families (cover-free families 20 are described in "Families of Finite Sets in Which No Set is Covered by the Union of Others," by P. Erdos, et al., Israel Journal of Mathematics 51 (1985), pp. 75-89). BE schemes based on such families are described, for example in "Efficient Methods for Integrating Broadcast Encryption and Traceability" by E Gafni, et al., Advances in

Cryptology - Crypto '99, Lecture Notes in Computer Science 1666 (1999), pp. 372-387

and "Coding Constructions for Blacklisting Problems without Computational Assumptions," by R. Kumar, et al., Advances in Cryptology - Crypto '99, Lecture Notes in Computer Science 1666 (1999), pp. 609-623. In particular, BE schemes that are
5 tight with the proven lower bounds on the total number of keys in such schemes are preferably employed herein.

In the preferred embodiment employing a simple randomized BE scheme, the key generator/manager module 15 generates, manages and continuously updates a universal set of broadcast encryption keys \mathbf{K} that are stored in database 16. For purposes of the
10 following discussion, the universe of keys \mathbf{K} (at any given time) in database 16 is denoted by $\mathbf{K} = \{k_1, \dots, k_K\}$, where subscript K denotes a total amount of keys in database 16. In the preferred embodiment, the set of keys \mathbf{K} in database 16 comprises a plurality of key subsets $U_1 \dots U_i$, wherein each key subset U_i is allocated to a corresponding one of the n subscribers $u_1 \dots u_n$. Each key subset U_i in database 16 comprises r keys that are
15 randomly selected from the set of keys \mathbf{K} and allocated to a given subscriber, i.e.,
 $(\forall i, |U_i| = r)$. The subset of keys associated with each subscriber is contained in a respective smartcard $S_1 \dots S_n$ which is issued to the subscriber by the service provider.

Furthermore, in the preferred embodiment employing a randomized BE scheme, the value of K is selected to ensure an (m, α) -cover free family with high probability.

20 More specifically, given the set of keys \mathbf{K} comprising a set of key subsets $U = \{U_1, \dots, U_n\}$, the set of key subsets comprises an (m, α) - cover-free family, if for all $i = 1, \dots, n$, and for all sets of m indices, $\{j_1, \dots, j_m\}$ not containing i , $|U_i \cap (\bigcup_{s=1}^m U_{js})| \leq (1 - \alpha)|U_i|$. A cover free family guarantees that it is impossible for any coalition of m cards to cover

an α -fraction of key contained on another smartcard. As described in greater detail below, the lower bound on the total number of active keys K (that are stored in database

16) to ensure an (m, α) -cover free family with high probability (at least $1-\epsilon$, where ϵ is

any positive fraction) is $\Omega\left(\frac{n(m+1)/(r(1-\alpha)+1)}{\epsilon}\right)$. A detailed proof for this lower

5 bound value for K is provided below.

It is to be understood that the factor of $n(m+1)/(r(1-\alpha)+1)$ in the bound on K is due to the *cover-free* requirement, that with very high probability it is impossible for m smartcards to cover another smartcard (although the construction itself remains viable for any value of K). Consequently, the BE scheme may require a very large key set K ,

10 where the total number of keys K is $\Omega(n^{m/r})$, when $r \geq m$, and $\Omega(n)$, otherwise (see, E.

Gafni, et al., “*Efficient Methods for Integrating Broadcast Encryption and Traceability*,” Advances in Cryptology - Crypto ‘99, Lecture Notes in Computer Science 1666 (1999), pp. 372-387). Indeed, these bounds may well be prohibitive for large user populations.

15 It is to be appreciated, however, that a long-lived BE scheme according to the present invention remains viable for values of K that fall below the lower bound value.

In other words, the *cover-free* requirement is not necessary for implementing a long-lived BE scheme according to the present invention. Indeed, a *cover-free* scheme gives no guarantee on the system’s behavior after $m+1$ cards are compromised. In addition, if a randomized attack model is considered (in which pirates pry open randomly selected 20 cards), it is demonstrated (below) that even in a system with significantly fewer keys than an m -*cover-free* system, a set of m compromised cards will cover another card only with negligible probability. Therefore, in the long-term analysis (and in the experimental

results) described below, it is to be appreciated that values of m are not defined and the *cover-free* requirement is not adhered to for such values. In other words, a long-lived BE scheme according to the present invention guarantees that a coalition of m cards chosen uniformly at random does not cover any other card (with very high probability) using 5 substantially fewer keys as required by a *cover-free* scheme.

Referring again to Fig. 1, to broadcast encrypted content, the broadcast encryption module 14 preferably employs a broadcasting protocol based on a secret sharing scheme that allows a privileged set of users P to decrypt an encrypted broadcast transmitted from the service provider using the keys contained on their smartcards. In a preferred

10 embodiment, as indicated above, the broadcasting protocol is an $(s, |S_P|)$ - threshold protocol which is used to broadcast a message M to a privileged set of users

$P = \{u_1 \dots u_{n-m}\}$, where S_P denotes a set of keys used to broadcast to the privileged set P , and where s ($s=ar$) denotes the number of keys out of the set U_i of r keys that a user needs to decode encrypted broadcast content. More specifically, with an $(s, |S_P|)$ -

15 threshold protocol being used to broadcast to privileged users P , K shares of M (M_{k1} , M_{k2}, \dots, M_{kK}) are created in such a manner that any s of the shares suffices to recover M .

The shares corresponding to key held by user in $X = \{u_{n-m+1}, \dots, u_n\}$ are discarded, and each remaining share is encrypted with its corresponding key and these encrypted messages are broadcast to the universe of users. In a preferred embodiment, $s = 1$, which

20 broadcasting protocol is sometimes referred to as an *OR* protocol as is known in the art.

A detailed discussion of threshold protocols can be found, for example, in the article by R. Kumar, et al, entitled: "Coding Constructions for Blacklisting Problems Without

Computational Assumptions," Advances in Cryptology – Crypto '99, Lecture Notes in Computer Science 1666 (1999), pp. 609-623 (hereinafter referred to as "*Kumar et al.*").

Threshold protocols are preferably employed herein as a component of a long-lived BE scheme because of their simplicity and because they yield broadcast encryption schemes with maximal resilience. A scheme is said to be ***m - resilient*** if ***m*** excluded (i.e., not privileged) users cannot recover the content even by pooling their keys. A broadcast encryption scheme with $(s, |S_P|)$ - threshold protocols for every privileged set P , is ***m-resilient***. In addition, the threshold protocols can be efficiently implemented with an error-correcting codes technique as described in *Kumar et al.* This technique involves generating shares of the content with an error correcting code. Consequently, the shares are much smaller than with a typical threshold secret sharing scheme, which results in a much smaller overhead (see D. Stinson, "Cryptography: Theory and Practice," CRC Press, 1995, for a detailed discussion on secret sharing).

Generally, the number of transmissions, t , in a broadcast encryption scheme is the number of keys that are used when encrypting the content. In the preferred embodiment in which the broadcast encryption module 14 implements an $(s, |S_P|)$ - threshold protocol to broadcast to a privileged set P , the number of transmissions is $|S_P| = t$.

When using threshold protocols for broadcasting, a user should have sufficiently many keys left after the keys of m other users are excluded to recover the content from the broadcast. Conventionally, this was achieved by allocating keys to users in such a way that the set system is a cover-free family. As indicated above, however, with the long-lived perspective, the cover-free requirement is not a necessary requirement for a long-lived BE scheme according to the present invention. A long-lived BE scheme

according to the present invention focuses on how the broadcast encryption scheme is affected by pirate smartcards, which are assumed to be smartcards containing at least r keys, and other unavailable cards (i.e., cards of excluded users). As indicated above, a card may be compromised either because of piracy or simple because a user ceases to be
5 an active subscriber and leaves the system. In either case, the keys on the card become permanently “unavailable” for use as encryption keys. A compromised (pirate) card may be a clone of some user’s card or may contain a set of r keys that does not exactly match any of the set of keys of the n users in the system.

Furthermore, it is important to note that a key becomes “unavailable” for
10 encryption use for one of two reasons: (1) the key is on a card belonging to an excluded user; or (2) the key is on a pirate decoder. The keys contained on a compromised smartcard (i.e., a pirate card or the card of an excluded user) are deemed “dead” (the compromised card is also deemed “dead”). On the other hand, a key that is not contained on a compromised card or one of an excluded user is considered “active.” A
15 smartcard is considered “clean” if it contains only “active” keys. The number of “dead” cards (i.e., either due to piracy or exclusion) is denoted d . It is to be understood that the reason behind the unavailability of a key has an effect on the BE scheme. When a key is dead because it appears on a compromised card, the key is permanently unavailable; whereas if it simple appears on an excluded user’s card, its unavailability may be short
20 term as the excluded user may be a privileged user at a later time.

In general, a long-lived BE method according to one aspect of the present invention for securely broadcasting to privileged users as smartcards become compromised comprises two basic components: (1) adjusting the set S_P of keys that are

used to encrypt the broadcast (e.g., by excluding “dead” keys from the set S_P and generating new keys to replace the dead keys) and (2) re-issuing smartcards to users (as necessary) whenever the number of “dead” cards d reaches a predetermined threshold.

The long-lived BE method is reactive in the sense that actions are taken in response to the

5 number of compromised smartcards (which, for example, might hamper the continuity of service for privileged users, or bring transmission costs to unacceptable levels). In the preferred embodiment, the expected number of “dead” cards that can be tolerated in the randomized recording scheme before it is necessary to record a user is at least

$$(\frac{K - r}{r}) \ln(1/\alpha), \text{ or approximately } \frac{K}{r} \text{ dead smartcards (a detailed proof of which is}$$

10 provided below).

In this regard, the long-lived BE scheme of the present invention is divided into “epochs,” where each *epoch* represents a period in which the number of d smartcards has exceeded the predetermined threshold. At the end of *epoch i*, $i = 1, 2, \dots$, a determination is made as to which smartcards, if any, need replacing, and new cards are issued to the

15 corresponding users. For the long-term analysis, a variable C_i^j denotes the set of smartcards in *epoch j* that were created in *epoch i* (note that $C_1^1 = \{U_1 \dots U_n\}$). The following table summarizes the notation used herein.

Table 1

<ul style="list-style-type: none">• $\{u_1, \dots, u_n\}$ is the set of all users.• U_i is the set of keys held by u_i.• $K = \{k_1, \dots, k_K\}$ is the set of all keys.• S_P is the set of keys used to broadcast to privileged set P.• n is the total number of users.• K is the total number keys.• r the number of keys per user.• t is the number of transmissions. Note that $S_P \leq t$.
--

- m is the number of users who are excluded.
- d is the number of unavailable (dead) cards at a certain point in time.
- C_i^j is the set of cards in epoch j that were created in epoch i .

It is to be appreciated that as indicated above, instead of the preferred randomized BE scheme described herein, the present invention may employ any suitable conventional BE scheme, such as BE schemes based on a *cover-free* family. For instance, a 5 randomized bucket-based scheme (such as described in Kumar, et al.) or a deterministic BE scheme based on polynomials (such as described in Gafni, et al.) may be extended to a long-lived BE scheme according to the present invention. Both constructions yield (m, α) - cover-free families. In the short term analysis (considering the first epoch), all three schemes are similar in terms of efficiency. The three BE schemes have very 10 similar costs. Each yield (m, α) -cover-free families with high probability when the total number of keys is close to the optimal bound of $n^{m/r}$ keys (this bound is for $\alpha=1/r$; i.e., no m users cover all another user's keys). In addition, all three schemes can tolerate approximately $\frac{K}{r}$ dead cards before recording is necessary. Since the schemes are so 15 close in terms of efficiency and cost, the simple randomized BE scheme is preferably employed as the basis of a long-lived BE scheme according to the present invention since it is the most simple, as it is entirely random, and it has the advantage over the bucket-based scheme of featuring r and m as independent parameters.

Consequently, the preferred long-lived BE scheme is an extension of the simple randomized BE scheme. As indicated above, the extension is reactive based on card 20 compromises and recording is performed once every d dead cards. For simplicity, the $(m, 1/r)$ -cover-free family version of the scheme (OR protocols). We emphasize that this

analysis is best applied to a stable privileged set of users, or to the entire set of users when m is small. In either case we expect to only have to record a user when too many of their keys appear on compromised (i.e., permanently unavailable) cards.

The main cost associated with long-lived schemes is the number of cards that
5 must be replaced. As demonstrated below, in the preferred long-live BE scheme, given a positive fraction β (and n , d , and r), at most βn of the cards need to be replaced during any recording session, provided K and the number of sessions is sufficiently large. This is shown below by the experimental results.

Referring now to Fig. 2, a flow diagram illustrates long-lived BE method
10 according to one aspect of the present invention. As indicated above, the long-lived BE method is preferably an extension of an efficient randomized BE method. Initially, each subscriber u_i of the broadcast service is allocated a set of decryption keys U_i (step 20). More specifically, in accordance with the underlying randomized BE scheme, each set of keys U_i is generated by randomly selecting r keys out of a set K comprising K keys. As
15 indicated above, the value K may be selected to ensure that an (m, α) -cover free family is obtained having high probability (where, preferably, $\alpha = 1/r$, i.e., no m users cover all of another user's keys). It is to be noted that $1/r \leq \alpha \leq 1$. In a preferred embodiment, the set of keys U_i are encoded on smartcards, which are then issued to the respective subscribers.

20 Next, a set of keys S_P are selected (from the set of keys K) to broadcast to a privileged set of users P (step 21). As indicated above, a preferred broadcasting protocol is an $(\alpha r, |S_P|)$ -threshold protocol, wherein a share is created for each key in the set S_P according to a $(\alpha r, |S_P|)$ -threshold secret sharing scheme ($s = \alpha r = 1$ denotes the number of

keys out of r that are needed to decode the content) and each share is encrypted with its corresponding key. In a preferred embodiment, $s=1$.

Over the course of time, one or more of the decryption keys that have been issued to the subscribing users may become compromised. As indicated above, a key is deemed
5 compromised (dead) if it is contained on a pirate smartcard or a smartcard of a subscribing user whose contract with the service provider has expired or been terminated (i.e., an excluded user). In addition, each pirate smartcard or smartcard belonging to an excluded user is considered dead. Therefore, when a pirate smartcard is discovered, it is analyzed and all the keys it contains are identified and deemed dead. Likewise, every
10 key contained on an excluded user's smartcard are deemed dead.

Accordingly, over the course of time, as compromised smartcards (pirate smartcards or smartcards of excluded users) are identified (affirmative result in step 22), to render such compromised/excluded cards useless, every key contained on such cards are deemed dead and, consequently, excluded from the set of keys S_P which is used to
15 broadcast to privileged users (step 23), thereby creating a new set S_P^1 (and so on).

Thereafter, broadcasting to privileged users is performed via a ($\alpha r, |S_P^1|$)–threshold protocol (and so on). The identified pirate smartcards and smartcards belonging to an excluded user are deemed “dead” or “unavailable.”

After compromised smartcards have been identified and deemed “dead”, a
20 determination is made as to whether the amount of “dead” smartcards d (at the given time) has met a predefined threshold (step 24). As indicated above, in a preferred embodiment, the expected number of dead cards d that can be tolerated before regarding

is needed is at least $(\frac{K - r}{r}) \ln(1/\alpha)$. If the number of dead cards d has not met the predefined threshold (negative determination in step 24), then the latest modified set of encryption keys (e.g., S_P^l) will be used to encrypt the content broadcast to privileged users (step 25). Thereafter, additional compromised smartcards may be identified over the course of time, resulting in continued modification (i.e., exclusion of the compromised keys) of the current set of keys used to encrypt the broadcast content (steps 5 22 and 23 are repeated).

On the other hand, if it is determined that the amount of dead cards d has met the predefined threshold (affirmative determination in step 24), a recording session is 10 performed. As indicated above, in a preferred embodiment, a recording session is performed once for each “epoch”, wherein each *epoch* is a time period when the number of dead cards d meets the predefined threshold. In a preferred embodiment, a recording session comprises updating the set of keys K by generating new keys to replace the dead keys (step 26). More specifically, let Z be the set of dead keys (i.e., all the keys 15 contained on the identified pirate and excluded smartcards) and let $z = |Z|$; note that $z \leq dr$. The update step comprises discarding all the keys in Z , and generating z new keys so that the resulting number of keys in the universal key set K is again K in total. The new set of all keys becomes $K' = (K \setminus Z) \cup Z'$. In short, new keys are generated for all 20 dead keys, and active keys are unchanged. The total number of active keys (i.e., keys that need to be stored by the broadcasting center (in database 16)) is unaffected by the recording procedure.

The next step in the preferred recording session is to determine if any (current) subscribers have less than a predetermined number of active keys on their smartcards (step 27). More specifically, in the preferred embodiment, a determination will be made as to whether there are any users having a current smartcard that contains less than αr active keys (preferably, $s = \alpha r = 1$ and, therefore no active keys). Each user having a smartcard with less than the predefined number of active keys (e.g., 0 active keys) will be recorded. Recording comprises the steps of allocating a new set of r keys by randomly selecting such keys from the updated set \mathbf{K} of K keys (step 28) and encoding the new keys on a new smartcard, which is then mailed to the user.

As indicated above, a key is considered dead either because it belongs to an excluded user or is on a compromised card. In the former case, the key is unavailable on what may be a short term basis, as an excluded user may well be a privileged user at another time. Hence, it is to be appreciated that the long term analysis described herein is best applied to a stable privileged set P , or to the whole set of users when the number of excluded users, m , is small. Given this, it is very likely that users will only be recorded when more than $(1 - \alpha)r$ of their keys are permanently unavailable (i.e., contained in compromised cards) rather than simply temporarily unavailable, due to the current set of excluded users. Indeed, the primary motivation for recording users is the presence of compromised (pirate) cards.

It is to be appreciated that the recording process (steps 26, 27 and 28) depicted in Fig. 2 is repeated for each epoch (i.e., for each affirmative result in step 24). As keys get compromised and smartcard are reissued to users, the users are partitioned into sets of users with cards with fresh keys, and users with cards containing keys some of which are

dead. This process is depicted in Figure 3. Recall that C_i^j represents the set of smartcards in epoch j that were created in epoch i . Initially in epoch 1,

$C_1^1 = \{U_1, \dots, U_n\}$. In a recording session in epoch 2, selecting (randomly) d dead cards

from C_1^1 yields C_2^2 , the set of users that need to be recorded, as well as $C_1^2 = C_1^1 \setminus C_2^2$; in

epoch 2, selecting d random cards from C_1^2 and C_2^2 yields C_3^3 as well as C_1^3 and C_2^3 ; and so on. As described in detail below, there comes a time in which steady state is reached where only a bounded number of users C_j^j need to be recorded in any epoch. A detailed proof of this steady state condition is provided below, along with an empirical analysis.

10 PROOFS

Lemma 1. As indicated above, the lower bound on the total number of active keys K (that are stored in database 16) in the randomized BE scheme to ensure an (m, α) -cover free family with high probability (at least $1-\epsilon$, where ϵ is any positive fraction) is

$\Omega\left(\frac{n^{(m+1)/(r(1-\alpha)+1)}}{\epsilon}\right)$. This lower bound value is derived as follows.

15 Consider $m+1$ users, u, u_1, u_2, \dots, u_m . First we bound the probability that u_1, \dots, u_m cover more than $(1-\alpha)r$ of u 's keys. Since $|\bigcup_{i=1}^m U_i| \leq mr$, we have the following bound:

$$\Pr(|U \cap (\bigcup_{i=1}^m U_i)| > (1-\alpha)r) \leq \frac{\binom{mr}{(1-\alpha)r+1} \binom{K-(1-\alpha)r-1}{ar-1} + \binom{mr}{(1-\alpha)r+2} \binom{K-(1-\alpha)r-2}{ar-2} + \dots + \binom{mr}{r}}{\binom{K}{r}}$$

20 Using binomial bounds and simplifying, we have:

$$\Pr(|U \cap (\bigcup_{i=1}^m U_i)| > (1 - \alpha)r) \leq \frac{(rem)^r e^{\alpha r} - 1}{K^{r(1-\alpha)} + 1}$$

Hence, the probability that $|U \cap (\bigcup_{i=1}^m U_i)| \leq \alpha r$ is at least $1 - \frac{(rem)^r e^{\alpha r} - 1}{K^{r(1-\alpha)} + 1}$.

There are $n-m$ privileged users, therefore the probability that there is at least one privileged user who shares more than $r(1-\alpha)$ keys with u_1, \dots, u_m , is at most

$$5 : 1 - \left[1 - \frac{(rem)^r e^{\alpha r} - 1}{K^{r(1-\alpha)} + 1} \right]^{n-m} . \text{ To account for all possible excluded sets of } m \text{ users,}$$

it suffices to multiply by $\binom{n}{m}$:

$$\left(\frac{n}{m} \right) \left(1 - \left[1 - \frac{(rem)^r e^{\alpha r} - 1}{K^{r(1-\alpha)} + 1} \right]^{n-m} \right)$$

Substituting a binomial approximation,

$$1 - \left(\frac{m}{ne} \right)^m \varepsilon \leq \left[1 - \frac{(rem)^r e^{\alpha r} - 1}{K^{r(1-\alpha)} + 1} \right]^{n-m}$$

10 If $K > r^{r+1} m^r e^{2r-1}$ (this is reasonable since we expect r to be small), then we can

use the fact that $(1-x)^{n-m} \geq 1 - (n-m)x$ when $x \leq 1$, to simplify this expression.

With this substitution, it suffices to show the following:

$$(n-m) \left[\frac{(rem)^r e^{\alpha r} - 1}{K^{r(1-\alpha)} + 1} \right] \leq \left(\frac{m}{ne} \right)^m \varepsilon$$

Solving for K yields the statement of *Lemma 1*.

Lemma 2: As indicated above, the expected number of dead cards d that can be tolerated in the randomized recording scheme before it is necessary to record a user is at least $\left(\frac{K-r}{r}\right) \ln(1/\alpha)$. This value of d is derived as follows.

The probability that a fixed key k is not held by a randomly chosen user is 5 $\left(1 - \frac{r}{K}\right)^d$. Therefore the expected number of keys that a user shares with d dead cards is, $r\left(1 - \left(1 - \frac{r}{K}\right)^d\right)$ (it is to be understood that because the long-lived BE scheme is completely random, we do not need to distinguish between cloned cards and non-cloned cards). We want to determine the values of d for which it is expected that recording is necessary:

$$10 \quad r\left(1 - \left(1 - \frac{r}{K}\right)^d\right) > (1-\alpha)r$$

$$d > \frac{\ln(1/\alpha)}{\ln\left(1 + \frac{r}{K-r}\right)}$$

using the fact that $\ln\left(1 + \frac{r}{K-r}\right) \leq \frac{r}{K-r}$, the statement of *Lemma 2* follows.

The next proof sets forth the bounds on the number of smartcards that need to be replaced in each *epoch*, namely, the (expected) size of set C_j^j in *epoch j*. Initially, we 15 first prove recurrence relations relating the expected number of cards in *epoch j* that were created in *epoch i* $i \leq j$, $E(|C_i^j|)$.

Lemma 3: In the randomized long-lived BE scheme with $\alpha = \frac{1}{r}$ and a fixed recording schedule of once every d dead cards, the following equations hold for all $i \geq 1$:

$$1. E(|C_{i+1}^{i+1}|) \leq \sum_{j=1}^i E(|C_j^i|) \left[1 - \left(1 - \frac{r}{K}\right)^{(i+1-j)d} \right]^r$$

$$2. \forall j, 1 \leq j \leq i, E(|C_j^{i+1}|) = E(|C_j^i|) \left(1 - \left[1 - \left(1 - \frac{r}{K}\right)^{(i+1-j)d} \right]^r \right)$$

5 The foregoing inequalities are derived as follows. To see the first inequality, note that if a user is recorded (or created) in epoch j , then during the time interval from the beginning of epoch j to the end of epoch i , $d(i+1-j)$ randomly chosen cards become unavailable. If these cards cover the user's card, then the user must be recorded. Due to the random

nature of the scheme, a user is covered with probability, $\left[1 - \left(1 - \frac{r}{K}\right)^{(i+1-j)d} \right]^r$.

10 We have a weak inequality rather than equality, because a user may be covered by fewer than $d(i-j+1)$ cards.

The second equation is obtained by noting that all users who were recorded (or created) in epoch j , and who are not covered by the end of epoch i , become the set of users C_j^{i+1} .

15 The inequalities from Lemma 3 are now used to perform an analysis that demonstrates that the long-lived BE scheme reaches a steady state in which only a bounded number of cards need to be replaced in each epoch. More specifically, the inequalities of Lemma 3 can be used to demonstrate (via approximate analysis) that an upper bound on the number of records per epoch holds in the limit, and that this upper

bound can be made small through appropriate choices of K, r and d . This analysis comprises three components. First, given fixed values of the parameters, there exists an integer ℓ_1 , such that the probability that a card is covered (and hence, needs to be refreshed) within ℓ_1 epochs, is negligible. The intuition for this result is that if a card has 5 been refreshed recently, then it is unlikely that it will be covered again within a small number of epochs. This result indicates that the contribution to $E(|C_{i+1}^{i+1}|)$ from the first ℓ_1 terms of inequality 1 of *Lemma 3* above is fairly small. Secondly, the later terms in inequality 1 may also not contribute much to the upper bound on $E(|C_{i+1}^{i+1}|)$. In particular, there exists an integer, ℓ_2 , (greater than ℓ_1), such that it is unlikely that a card 10 will not be covered within ℓ_2 epochs. Note that this implies that when $i - j \geq \ell_2$, $E(|C_j^i|)$ is fairly small, and hence, will not contribute much to the upper bounds on $E(|C_{i+1}^{i+1}|)$. Lastly, ℓ_1 and ℓ_2 are on the order of K/rd , hence the dominating terms are those for which $i-j$ is $\Theta(K/rd)$, and this leads to an approximation for the upper bound of the steady state record rate, β . The following proof demonstrates this analysis in 15 detail.

Lemma 4: Assume n, K, r, d and $\epsilon > 0$ are given, the following are true:

1. If $\ell_1 \in O\left(\frac{\epsilon \sqrt{rK}}{rd}\right)$, then the probability that a card is covered within ℓ_1

epochs is less than ϵ .

2. If $\ell_2 \in \Omega\left(\frac{(1-\varepsilon)^{1/r} K}{rd}\right)$, then the probability that a card survives for more than ℓ_2 epochs before it is covered is less than ε .

3. If $i-j \in \Theta\left(\frac{K}{rd}\right)$, then the coefficient of $E(|C_j^i|)$ in inequality 1 of the

above *Lemma 3* is approximately $\left(1 - \left(\frac{1 - r/K}{e^c}\right)^d\right)^r$, where c is a constant.

5 The proof is as follows:

1. The probability that a card is covered by $d\ell_1$ randomly chosen cards is $\left(1 - \left(1 - \frac{r}{K}\right)^{d\ell_1}\right)^r$. Setting this quantity less than ε and solving for ℓ_1 yields

$$\ell_1 < \frac{\ln(1 - \varepsilon^{1/r})}{d \ln(1 - r/K)}.$$

10 2. The probability that a card is not covered within ℓ_2 epochs is $\left(1 - \left(1 - \frac{r}{K}\right)^{d\ell_2}\right)^r$. Setting this quantity less than ε and solving for ℓ_2 yields

$$\ell_2 > \left(\frac{K}{rd}(1 - \varepsilon)^{1/r}\right)$$

3. Assuming that $i - j = \frac{cK}{rd}$ for some constant c , we'll bound the contribution of $i - j = \frac{cK}{rd}$ to the inequality in *Lemma 3* (i.e., we'll bound the coefficient

of $E(|C_j^i|)$ - given the earlier results, this bound is an approximate upper bound to β , the

long term steady state. When $i - j = \frac{cK}{rd}$, the coefficient of $E(|C_j^i|)$ is

$$\left(1 - \left(\left(1 - \frac{r}{K}\right)^d\right)^{\frac{cK}{rd} + 1}\right)^r. \text{ When } K \text{ is sufficiently large, this is the order of}$$

$$\left(1 - \frac{\left(1 - \frac{r}{K}\right)^d}{e^c}\right)^r.$$

When combined with parts 1 and 2, part 3 of this Lemma indicates that the steady state record rate β should decrease with K (which agrees with the experimental results that follow). The quantity also increases with d , which agrees with the basic intuition that the longer we wait to record, the more recording we will have to do.

Experimental Results

The following experiments illustrate the dynamics of smartcard replacement in accordance with a long-lived BE scheme according to the present invention and corroborate the steady state claim. In these experiments, equation 1 of Lemma 3 was evaluated (assuming an equality rather than an inequality) for a variety of parameter settings to track the number of cards that were issued in every epoch. Furthermore, we focused on the random attack model, and assumed that, in each epoch, the dead cards d are selected uniformly at random from the set of user cards. As indicated above, a long-lived BE scheme according to the present invention guarantees that a coalition of m cards chosen uniformly at random does not cover any other card (with very high probability) using substantially fewer keys. Thus, in the long-term analysis, and in the following simulation, we did not adhere to the *cover-free* requirement and, consequently, the value

of K was not constrained by the bound of *Lemma 1*. Instead we de-coupled the number of users n from the total number of keys K number, and observed the behavior of the resulting schemes in terms of how any cards need to be issued per epoch, and let K be a free parameter which we varied.

To justify this decoupling of K from n , we present the following lemma that provides a lower bound on K such that with high probability, none of the n user cards are covered by d randomly chosen compromised cards (i.e., some cover-freeness is achieved with high probability). As mentioned above, this lower bound may be much smaller than the size of K in a d -cover-free family.

Lemma 5: Assume $\epsilon > 0$, n , r and d are given. If $K > r/c$, where c is a constant that depends on ϵ , n , r and d , then the probability that any user's card is covered by d randomly chosen cards is less than ϵ . The proof is as follows.

The probability that n (randomly chosen) user's cards are not covered by d randomly chosen cards is $\left[1 - \left(1 - \left(1 - \frac{r}{K}\right)^d\right)^r\right]^n$. Hence, we solve the following inequality for K , $\left[1 - \left(1 - \left(1 - \frac{r}{K}\right)^d\right)^r\right]^n > 1 - \epsilon$, which yields

$$K > \frac{r}{\left(1 - \left(1 - \left[1 - (1 - \epsilon) \frac{1}{n}\right]^{\frac{1}{r}}\right)^{\frac{1}{d}}\right)}.$$

As an example, for the values of n , r and d used in Fig. 4 below, and $\epsilon = .1$, the above *Lemma 4* give a lower bound on K of approximately 69, which is far less than the lower bound of approximately 10^{10} for a d -cover-free family.

In each of the following experiments, the user population was set to a size $n = 100,000$, which is viewed as being on the low end of real population sizes. The card capacity r ranged between 10 and 50, which is realistic for current smartcards with 8KB of memory and keys requiring, say, 64 bytes each including overhead. The epoch length 5 was varied to be $10 \leq d \leq 50$ dead cards. The number of key was varied between $1000 \leq K \leq 5000$, which values are smaller than the number of keys required by Lemma 1, which calls for $K \geq n^{m/r}$ keys (note the dependency on m , the number of users the underlying BE is able to exclude)

Figure 4 is a diagram that illustrates the dynamics of the card reissue strategy, and 10 the effect of the total number of keys K . More specifically, Fig. 2 illustrates 3 curves each showing the number of cards that are reissued per epoch for $n=100,000$, $d=20$, for values of $K= 1000, 2500$ and 5000 , respectively. We see that each curve begins with oscillations. In the first epoch ($d=20$) of each curve, we see that no cards are re-issued (since the first dead cards do not cover any user). But after a certain number d_c of dead 15 cards are discovered, enough keys are compromised and there is a rapid increase in card reissue. This in turn “cleans” the card population and, as shown, the card reissue rate drops. We see that the oscillations are dampened and a steady state appears fairly quickly, especially for higher values of K .

Indeed, as is shown in Fig. 4, the parameter K affects several aspects of the 20 dynamics: the first card reissue point d_c of each curve occurs later as K increases ($d_c \approx 40$ for $K = 1000$ whereas $d_c \approx 200$ for $K = 5000$). In addition, as K increases, the curve oscillations are gentler, have a smaller amplitude and lower peak rate. Most importantly,

the steady state rate of card reissue is lower for larger K . For instance, steady state occurs at ≈ 9400 cards per *epoch* for $K=1000$, whereas steady state occurs at ≈ 2000 cards per *epoch* for $K=5000$. Overall, we see that increasing K improves all the aspects of the reissue strategy. Therefore, it is preferable to use the largest possible K that is within the 5 technological requirements.

Next, Fig. 5 is a diagram that illustrates the effect of increasing the card capacity r . More specifically, Fig. 5 illustrates three curves representing the number of cards reissued per epoch, with $n=100,000$, $K=5000$ and $d=20$, for values of $r = 30, 20$ and 10 . We can see that as the value of r increases, the cost of card reissue increases: the curves 10 having larger values for r result in higher steady state reissue rates and higher peak reissue rates. This is consistent with the fact that as r increases, we expect each key to be contained in more cards, so the effect of a compromised key is more widespread. In addition, as indicated above, we expect to have to record users sooner when r is large 15 (and K is fixed). Having a smaller r , however, means that the expected transmission length would be longer; at the extreme, setting $r = 1$ would give optimal reissue rates (no cards need to be reissued), with very long transmissions.

Next, Fig. 6 illustrates the effect of increasing the epoch length d . More specifically, Fig. 6 illustrates three curves representing the accumulated total number of reissued cards, with $n=100,000$, $K=5000$ and $r=10$, for epoch lengths $d=20, 40$ and 80 . 20 Fig. 6 illustrates that a longer epoch length results in a smaller total number of reissued cards. On the other hand, as the length of the epoch increases, the number of compromised keys during each epoch increases, thereby possibly resulting in many users

not being broadcasted to securely (in the absence of unicasts). Consequently, the recording costs and transmission costs may influence the choice of d .

In summary, a long-lived broadcast encryption method according to the present invention provides a mechanism for adapting broadcast encryption to the presence of pirate decoders and maintaining the security of broadcasts to privileged users while rendering all pirate decoders useless. A long-lived broadcast encryption scheme is a more comprehensive solution to piracy than conventional methods such as traitor tracing schemes, because the latter only seek to identify the makers of pirate decoders and do not provide a mechanism for maintaining secured broadcasts once keys have been compromised. Indeed, one fundamental difference between a long-lived BE scheme as described herein and conventional techniques is that the long-live BE scheme seeks to minimize the amount of communication (recording) that is needed by adapting to the presence of compromised cards (or equivalently, revoked users) by removing the keys contained on a pirate decoder (or revoked user's cards) from the encryption process.

Consequently, a long-lived BE scheme as described herein is a more efficient long-term solution than conventional revocation schemes because is minimized the amount of recording that must be done in the long-term.

In addition, as described above, although a conventional broadcast encryption scheme may only be designed to allow any set of m users to be excluded, a long-lived BE scheme according to the present invention can tolerate a high number of compromised cards (in addition to the m excluded cards) before any users need to be recorded. Indeed, through an appropriate choice of parameters, the affected number of users (who need to be recorded in a given epoch during steady state) can be a small fraction of the total

number of users. That is, for a long-lived broadcast encryption scheme according to the present invention, for any positive fraction β , the parameter values may be selected, *a priori*, in such a way to ensure that eventually, at most β of the n cards must be replaced in each epoch.

5 Although illustrative embodiments have been described herein with reference to the accompanying drawings, it is to be understood that the present system and method is not limited to those precise embodiments, and that various other changes and modifications may be affected therein by one skilled in the art without departing from the scope and spirit of the invention. It is to be understood that all such changes and
10 modifications are intended to be included within the scope of the invention as defined by the appended claims.